

LSI 01-642

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR PATENT

ON

*A MECHANISM FOR ENCODING AND DECODING UPGRADEABLE RPC/XDR  
STRUCTURES*

BY

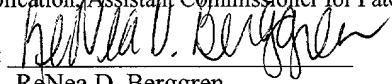
WILLIAM P. DELANEY  
14715 SIEFKES STREET  
WICHITA, KS 67230  
CITIZEN OF USA

CERTIFICATE OF MAILING BY "EXPRESS MAIL"

"Express Mail" Mailing Label Number EV 013 245 174 US

Date of Deposit: November 19, 2001

I hereby certify that this correspondence is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. § 1.10 on the date indicated above and is addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231

BY:   
ReNea D. Berggren

# *A MECHANISM FOR ENCODING AND DECODING UPGRADEABLE RPC/XDR STRUCTURES*

## FIELD OF THE INVENTION

5 [0001] The present invention generally relates to the field of computer network communications, and particularly to a method of communication involving encoding and decoding data between devices having dissimilar structure definitions.

## BACKGROUND OF THE INVENTION

10 [0002] In a distributed system, the data is processed by a set of components connected by a network. The components may have different hardware and/or operating systems and may send and receive data using different transfer syntaxes. Data often needs to be in a form which permits the data to be distributed across the network by converting the data into the format required for the destination component's hardware or operating system.

15 [0003] The Remote Procedure Call (RPC) package defines an industry-standard procedure calling model for distributed applications by invoking an independent set of functions used for accessing remote nodes on a network. The RPC model is derived from the programming model of local procedure calls. An RPC executes a procedure located in a separate address space from the calling code. It uses a procedure declaration for every procedure call. All calls to a procedure must conform to the procedure declaration.

20 The RPC protocols extend the concept of local procedure calls across the network, which means distributed applications may be developed for transparent execution across a network.

25 [0004] The External Data Representation (XDR) defines a standard representation for data in the network to support heterogeneous network computing. The XDR standard data representation convention is a set of library routines that allows a programmer to describe arbitrary data structures in a machine-independent fashion. By using XDR, systems do not have to understand and translate every data format that may exist on the network as there is only the one convention. Data is translated into XDR format before it

is sent over the network and, when received, is translated into the data convention used there. New computer architectures can be incorporated into the network without requiring the updating of translation routines.

[0005] The current technique for specifying the information exchanged between RPC clients and servers is based on the use of an interface definition file that uses syntax similar to C code. This file contains a specification of each data element that can be exchanged. The type of each such element can be either a primitive (built-in) type, such as integer, character, Boolean, etc., or a complex (derived) type, such as a structure or union. For RPC applications with even modest complexity, the vast majority of exchanged data elements are of a complex type, with most common being structure and union.

[0006] To facilitate the development of client and server application software, the RPC/XDR model uses a tool that automatically generates code to translate between the network representation of a data element and the language-specific implementation of that element. Since structure data types are the most relevant for this environment, techniques used for generating translation code for them are the point of focus. A structure is really just an ordered collection of fields, where each field has its own characteristic data type. Consequently, in the RPC/ XDR model, the code to translate a structure consists of a sequence of statements, each of which translates one field within the structure. If a field is of a primitive (built-in) data type, there are built-in library routines within the RPC/ XDR package that handle the translation, so the code generator creates code that calls into the library. If a field is of a complex type, the code generator creates code that translates it by invoking the auto-generated code for that field's complex type.

25 [0007] The key restriction with the current methodology is that the side that is encoding the data and the side that is decoding it must be using auto-generated code that was derived from exactly the same version of the interface definition. The reason for this is that the encoder simply dumps the encoded version of each successive field into the message stream, and the decoder extracts the exact same number of fields from the

message stream at the other side of the connection. If the decoder were to be based on an interface definition for the structure that contained more fields than the encoder knew about, the decoder would try to extract data for the additional fields from the message stream, even though the encoder did not place any such data in the message. On the other 5 hand, if the decoder's structure definition contained fewer fields than the encoder's, the decoder would finish decoding and leave extra data in the message stream. This would interfere with further decoding of structures and other data elements that follow the structure in question, since the decoders for those ensuing elements would try to decode data that is really not part of their elements.

10 [0008] The standard resolution for this problem in the RPC/ XDR environment is to simply create new structures, in addition to the old ones, and modify all the relevant elements of the interface definition to ensure that both the old and the new/ enhanced structures are used where appropriate. This generally requires the development of an entire new interface definition, and the need for new code to support both the new and the 15 old versions of the interface. Over time, this evolves into a need to support multiple (i.e., even more than two) versions of the interface, which is a serious maintenance issue.

[0009] The only real alternative is to use a lock-step migration strategy. When changes are needed to the interface definition, they are made on both the client and the server. All clients and all servers must then be upgraded to the new interface definition in a lock-step 20 deployment. This approach reduces the amount of code required, and thus the maintenance effort, but imposes extreme difficulties in the deployment phase, since incompatible versions of the interface must not coexist.

[0010] Therefore, it would be desirable to provide a method and system for communicating between network elements having dissimilar data structure definitions.

25

#### SUMMARY OF THE INVENTION

[0011] Accordingly, the present invention is directed to a method and system for data communications between network elements in a heterogeneous network which is platform independent.

[0012] This invention is an extension of current techniques for automatically generating language-specific translation code for information that flows between heterogeneous systems via network interactions. In particular, this invention focuses on a new technique that can be applied to the Remote Procedure Call (RPC) and External Data 5 Representation (XDR) frameworks standardized by Sun Microsystems. With the current approach, the client and server side of an RPC connection must operate with exactly the same structure definitions for all data objects that are passed from one side to the other. This requirement stems from the fact that the automatically-generated code that translates 10 between the network and local data representations has no provisions for managing situations where one side is using an up-level or down-level version of the object definitions relative to the other side. This invention provides for extensions to the existing methodology by which the translation code is generated so that the client and server can correctly operate even when they are using different versions of the data objects identified in the RPC interface definition. It is described specifically as a solution 15 to the problem of versioned structures in the RPC/ XDR environment, but the essence of the invention is applicable in any environment where information needs to be encoded/ decoded for exchange between heterogeneous entities.

[0013] This invention relates to a system and method for communicating across a heterogeneous network having components with dissimilar data structure definitions, 20 comprising: 1) prefixing all encoded data structures with a length value that reflects the size of the encoded data structure; 2) for an up-level sender, if the up-level definition of data elements is greater than the down-level definition of data elements, the down-level receiver reads the length value and decodes the encoded data structure according to the receiver's data definition and upon completion of decoding, the receiver determining the 25 amount of the encoded data structure that was decoded and skipping the remainder of the encoded data structure according to the length value; 3) for a down-level sender, if the up-level definition of data elements is greater than the down-level definition of data elements, for built-in type data fields, automatically assigning a default value to any field for which the received data has provided no value, and, for derived type data fields,

calling an initialization routine which assigns a default value to any built-in type data field or calls the initialization routine for a derived type data field; and 4) to avoid the complexities associated with assignment of default values and initialization routines following a set of predetermined rules which include extending only data structures  
5 which are passed from a server to a client; ensuring that a down-level client's ignorance of extended data causes no ill effects in that client's operational behavior; allowing down-level clients to interact with up-level servers and disallowing up-level clients from interacting with down-level servers; and, in cases where extensions are needed for data structures passed from a client to a server, defining a new data structure that includes  
10 both old data fields and new data fields.

15 [0014] It is to be understood that both the forgoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention as claimed. The accompanying drawings, which are incorporated in and constitute a part of the specification, illustrate an embodiment of the invention and together with the general description, serve to explain the principles of the invention.

#### BRIEF DESCRIPTION OF THE DRAWINGS

20 [0015] The numerous advantages of the present invention may be better understood by those skilled in the art by reference to the accompanying figures in which:  
[0016] FIG. 1 illustrates both the traditional XDR encoding scheme and the enhanced encoding scheme for two different exemplary structures; and  
[0017] FIG. 2 illustrates a table of the different modes of operation.

#### DETAILED DESCRIPTION OF THE INVENTION

25 [0018] Reference will now be made in detail to the presently preferred embodiments of the invention, examples of which are illustrated in the accompanying drawings.  
[0019] Referring generally now to FIGs. 1 and 2, an exemplary embodiment of the present invention is shown.

[0020] It is common to interconnect a variety of computer architectures and computer operating systems in a single network. In a heterogeneous distributed computing environment there will be differences in the way that data is represented on different machine architectures. To handle these differences External Data Representation (XDR) 5 library routines may be used to encode all components of a message before a transmission, and to decode the components of a message after reception.

[0021] One configuration commonly used for performing operations over a network is a client/server architecture. A heterogeneous network may include various clients and servers such as a Windows client, a Windows NT server, and a Unix server and may 10 include software applications, operating systems, hardware products, peripheral devices, and system hardware. Servers provide services and include database servers, transaction servers, groupware servers, file servers and object servers. A client process that is executing either on a server computer or another computer is a consumer of services provided by the server. Each interaction between a client and a server tells a server 15 which service is requested. After the server receives the request, the server determines how to perform the service specified in the request.

[0022] Communications between a client and a server over heterogeneous network require a method for transporting requests over network from a client running under one operating system to a server that is either running under another operating system, or the 20 same operating system. After the server component management function module determines received data in a format compatible with a server computer, the server component management function module reads the version specified in the received data. One widely used method for communication over heterogeneous network is a remote procedure call (RPC).

[0023] The use of RPCs introduces limitations on updates and modifications. The 25 standard resolution for this problem in the RPC/ XDR environment is to simply create new structures, in addition to the old ones, and modify all the relevant elements of the interface definition to ensure that both the old and the new/ enhanced structures are used where appropriate. This generally requires the development of an entire new interface

definition, and the need for new code to support both the new and the old versions of the interface. Over time, this evolves into a need to support multiple (i.e., even more than two) versions of the interface, which is a serious maintenance issue. Alternatively, when changes are needed to the interface definition, they are made on both the client and the 5 server. All clients and all servers must then be upgraded to the new interface definition in a lock-step deployment to support each version, regenerated, recompiled, and relinked for each of the computers on the network. For heterogeneous networks, this is a formidable task. Thus, while the trend is to implement RPC for component management over a heterogeneous network, the requirement of the current RPC architecture for either 10 consistent versions or support of all versions throughout such a network will limit the actual utilization of RPC for component management.

15 [0024] With the enhanced technique of this invention, it is possible to upgrade the data structure definitions within a given RPC/ XDR specification, such that clients and servers at different versions may coexist. Furthermore, the new technique does not require the server to implement a multitude of versions of the interface definition.

20 [0025] In the enhanced methodology, the sender of a given data element always encodes it using the version of the data definition that is “current” for the sender. This statement applies without regard to whether the sender is the client or the server. The receiver, on the other hand, must be able to accommodate receipt of data elements that were encoded using either newer (i.e., up-level) or older (i.e., down-level) versions of the interface 25 definition. Equivalently, the receiver must handle data elements in the in-bound message that are either larger or smaller than expected, given the receiver’s operating definition of the data elements. The techniques for handling these two cases are presented below.

[0026] A situation in which messages with more-than-expected data elements can be received if the sender is up-level relative to the receiver, and the up-level definition of the structure contains more data elements than the down-level definition being used by the receiver. It can also occur if the sender is down-level and the up-level definition being used by the receiver contains fewer data elements. The latter case is more difficult to

handle, so a reasonable restriction is imposed that up-level definitions must always extend the data structures with additional elements.

[0027] To handle the former case, the RPC/ XDR data conventions are modified as follows. First, when encoding a structure, the sender will prefix the encoded version of the data with a length value that reflects the encoded size of the structure, as shown in FIG. 1. (All encoded data has a length value.) Second, when decoding a structure, the receiver will read the length prefix, then decode the data according to the receiver's data definition. Upon completion of this process, the receiver determines the amount of encoded data that was consumed. If this amount is less than the known encoded length of the structure, the extra amount is simply skipped. This has the effect of positioning the receiver's decoding algorithm at the appropriate offset in the message to begin decoding the next data element. For example, in the structure XXX in FIG. 1, the sender uses an up-level definition that includes some new field named c (following the existing fields a and b). In this case, the encoded length value (len1) will include the encoded length of field c.

However, the decoder only knows how to decode fields a and b. Upon completing this known work, the decoder will detect that some amount of un-decoded data remains for structure XXX, and will simply skip over it. This applies even when the decoder is operating on an embedded version of structure XXX as a field of structure YYY.

[0028] A situation in which messages with fewer-than-expected data elements can be received is if the sender is down-level relative to the receiver, and the down-level definition of the structure contains fewer data elements than the up-level definition being used by the receiver, see FIG. 2. It can also occur if the sender is up-level, and the down-level definition being used by the receiver contains more data elements. As indicated above, the latter case imposes extra difficulties, and is thus restricted. To handle the former case, the RPC/ XDR data conventions are modified as follows. First, as stated above, the sender must prefix the encoded version of the data with a length value that reflects the encoded size of the structure. Second, when decoding a structure, the receiver will first read the length prefix. In this situation, there may be some fields of the

receiver's structures that have no encoded value in the data from the sender. Consequently, the decoding algorithm used on the receiver side will always check, before beginning to extract data for each field of the structure, to see if the encoded length has already been reached. If it has not, decoding proceeds in the usual way. However, if all  
5 of the encoded data has been consumed, the decoding routine will set the receiver's additional fields to a "default" value. The specification of default values is handled via extensions to the interface definition language used for RPC/ XDR.

[0029] The only real extensions required of the RPC/ XDR interface definition language are for support of default value specifications for structure fields that are not present in a  
10 received message. A given field can be either a built-in type or a complex (derived) type.

[0030] For built-in types, an appropriate "default specification" in the RPC/ XDR language is simply a grammar extension of the form "= default-value", where default-value is an appropriate value for that type (i.e., integer, floating point, Boolean, string, etc.). This extension allows automatic generation of code for decoding data that will  
15 assign the appropriate value to any field of a built-in type that is not present in the inbound message.

[0031] For fields of a derived type, the problem is a bit more difficult. However, it can be solved using the same basic approach as above, where only built-in types have a default value specification. The code generator of XDR translation routines creates an  
20 appropriate initialization routine for each defined derived type. This initialization routine contains code that sets each field within the derived type to its default value by either directly initializing the value for a field of built-in type, or by calling the initialization routine for a field of a derived type.

[0032] From the preceding discussion, it is clear that substantially more effort is required  
25 to handle the case where received data contains fewer-than-expected data elements. This case requires changes to the encoding and decoding logic, and non-trivial extensions to the interface definition language. On the other hand, the case where more-than-expected data elements are received can be handled with less overall effort, since none of the

interface definition changes are required. It is possible to exploit just this subset of the enhancement, provided the following rules are followed:

[0033] 1) Upgrade (i.e., extend) only those structures that are passed from the server to the client. Ensure that a down-level client's ignorance of the extended data causes no ill effects in that client's operational behavior. This is done through the careful extension of the structures. It depends heavily on the specific details of the interface being modified.

[0034] 2) Allow down-level clients to interact with up-level servers, but disallow up-level clients from interacting with down-level servers.

[0035] 3) In cases where a new version (i.e., up-level) of the interface requires modifications/ extensions are needed for structures passed from client to server, it will be necessary to define a new structure that includes both the old fields and the new fields. That is, the upgradeable structure support must not be used in these cases, since it would violate the rules that allow just the subset of this enhancement to be used (this corresponds to the case of FIG. 2 for an up-level sender in which the up-level structure has more data elements).

[0036] For this subset, the interface definition language extensions can be omitted, resulting in a simpler implementation.

[0037] It is believed that the A MECHANISM FOR ENCODING AND DECODING UPGRADEABLE RPC/ XDR STRUCTURES of the present invention and many of its attendant advantages will be understood by the forgoing description. It is also believed that it will be apparent that various changes may be made in the form, construction and arrangement of the components thereof without departing from the scope and spirit of the invention or without sacrificing all of its material advantages. The form herein before described being merely an explanatory embodiment thereof. It is the intention of the following claims to encompass and include such changes.